

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization International Bureau



(43) International Publication Date
3 June 2004 (03.06.2004)

PCT

(10) International Publication Number
WO 2004/047375 A1

(51) International Patent Classification⁷: H04L 12/44,
29/12, 12/64, 12/56, 12/46

[SG/SG]; Blk 213 Bedok South Avenue 1 #02-02, Singapore 469337 (SG).

(21) International Application Number:
PCT/SG2002/000268

(74) Agent: WATKIN, Timothy, Lawrence, Harvey, Lloyd Wise, Tanjong Pagar, P O Box 636, Singapore 910816 (GB).

(22) International Filing Date:
15 November 2002 (15.11.2002)

(81) Designated States (national): CN, KR, US.

(25) Filing Language: English

(84) Designated States (regional): European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR).

(71) Applicant (for all designated States except US): INFINEON TECHNOLOGIES AG [DE/DE]; St.-Martin-Strasse 53, 81669 Munich (DE).

Declaration under Rule 4.17:


— of inventorship (Rule 4.17(iv)) for US only

Published:

— with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: REDUCING THE MEMORY REQUIREMENTS OF A DATA SWITCH



(57) Abstract: A data switch (4) is proposed which can be operated such that it generates a MA table of MAC addresses of devices associated with some of its ports (9, 13, 15, 17, 19), but not the MAC addresses associated with one of the ports (11). Packets received by the switch, other than from the port (11), are assumed to be destined for the port (11) and are transmitted via the port (11). If the port (11) is connected to a network, and all the other ports are connected to a relatively small number of devices, such as one device per port, then the memory requirement to store the MA table is dramatically reduced.

Reducing the memory requirements of a data switch

Field of the invention

The present invention relates to a data switch. The switch is particularly suitable for use in Ethernet voice communication devices.

5 Background of Invention

Data switches, such as Ethernet switches, are well known which transfer data packets between ingress/egress ports connected by a switching fabric controlled by a control unit. Two or more of the ports of the switch are connected (either directly or via other devices or networks) to devices having 10 MAC (medium access control) addresses. Each of the data packets contains a destination address specifying the MAC address of the device to which the data packet is directed and an origin address specifying the MAC address of the device where the packet originated.

When the switch is first turned on, it typically does not "know" the MAC 15 addresses of the devices connected to it, and which ports those devices are connected to. Typically, the device learns this information automatically using the data packets it receives, by extracting the origin addresses from the data packets and noting which port the data packet arrived at. The association between these two pieces of information is stored in a MAC address table 20 ("MA table"). A variety of algorithms are known for constructing the table. When a data packet is received containing a destination address, the switch can look up the address in the MA table, find the associated port, and transmit the data packet to that port, for onward transmission to the device having that MAC address.

25 A significant problem, however, is that if the switch is transmitting data among a very large number of devices, the size of the MA table must become very great. This is especially true because many of the algorithms for generating

MA tables are inefficient in terms of the amount of writable memory required in relation to the amount of data which it actually has to store.

Summary of the Invention

5

The present invention aims to provide a new and useful data switch, and in particular one with a lower memory requirement.

In general terms, the invention proposes that the MA table is not constructed for MAC addresses associated with at least one of the ports of the switch, 10 called here a "first port" of the switch. When the switch identifies that a certain MAC address is associated with this first port, it does not insert this information into the MA table.

The invention is particularly applicable to switches which are intended to be used with a single one of the ports connected directly to a communications 15 network, such as a LAN, and with the other ports connected directly to devices (or unused). In such circumstances, the first port is preferably the one connected to the LAN, so that the data switch may learn the MAC addresses associated with the devices, but not the MAC addresses associated with the port which is connected to the network. This means that the device will be 20 capable of correctly transferring data packets to the ports connected to the devices.

When an packet which is received from a port other than the first port, and the destination address is not one of the learnt MA addresses, the switch may assume that the destination address is one of the (unlearnt) MAC addresses 25 associated with the first port, and so transmit the packet to the first port for transmission to the network. This concept provides an alternative expression of the invention.

The first port of the switch may be predetermined (e.g. a port which is intended always to be the one connected to a LAN). Alternatively, the first port may be determined by a control signal applied from outside the switch, or automatically, e.g. based on an automatic determination of which data port

5 receives the greatest number of new packets, or which data port is associated with the greatest number of MAC addresses.

Preferably, the switch is provided with a first control register for storing a control bit, and the switch is arranged, based on the value of the control bit, to either learn MA addresses selectively according to the invention, or according

10 to conventional learning mechanisms (i.e. with MAC addresses for all ports being learnt). This control bit may be set by a control signal from outside the switch, and this allows the switch of the invention to be operated as a conventional switch when desired.

In a preferred embodiment, the switch is arranged to have at least one port for

15 receiving/transmitting voice signals, i.e. data packets containing voice data collected by a microphone and a destination address indicating where the data packets are to be sent.

For example, it is known to provide a voice communication device for supporting communication with a second device. The communication device

20 includes a microphone, a speaker, circuitry for transforming sound signals received from the microphone into data packets and for transforming data packets into control signals for the speaker, and a data switch having a port connected to the circuitry. The data switch is connected to the second device via a network such as a LAN network. Typically, the data switch includes one
25 or more further ports for connection to other devices, such as personal computers. Such a system may be implemented using a switch according to the invention. In this case, the first port of the switch is preferably the one connected to the network.

In this embodiment, the switch will typically have a small number of other ports, e.g. no more than 5, no more than 10 or no more than 15, and these other ports will be associated with a relatively low number of MAC addresses, typically no more than one or two. Thus, the maximum size of the MA table
5 required to learn the MAC addresses associated with all the ports apart from the first port will typically only have to be enough to learn no more than 5 to 30 data items.

Brief Description of The Figures

Preferred features of the invention will now be described, for the sake of
10 illustration only, with reference to the following figures in which:

Fig. 1 is a schematic illustration of an embodiment of the invention.

Detailed Description of the embodiments

15 Referring to Fig. 1, the embodiment illustrated is a voice communication device 1 including a microphone 3 and a speaker 5. The voice communication device 1 includes a socket 2 by which the voice communication device 1 can be connected to a LAN (not shown) by means of a plug 6 attached to a cable 8 leading to the LAN. The LAN includes a number of devices, each having a
20 respective MAC address. At least one of these devices is another voice communication device, and the voice communication device 1 is for communication with a selected one or more of these other voice communication devices.

25 The microphone 3 and speaker 5 are connected to circuitry 7 which converts sound signals captured by the microphone 3 into data packets including the destination address of the other voice communication device(s) with which the voice communication device 1 is in communication. The circuitry 7 is also adapted to convert data packets originating in the other voice communication

device(s) into signals for transmission to the speaker 5 for conversion into sound signals.

The voice communication device includes a data switch 4, having ports 9, 11, 5 13, 15, 17, 19. Each of these is an ingress/egress port. The port 9 is connected to the circuitry 7, and the port 11 is connected to the socket 2 for connection to the LAN network. The ports 13, 15, 17, 19 are connected to respective sockets 23, 25, 27, 29, to which other devices, such as personal computers may be connected. Although four such sockets are shown, the 10 number of sockets is not limited in this respect.

The switch 4 includes a switching fabric (not shown, but of a conventional design), and a control unit (not shown) for controlling the switching fabric. The control unit includes a first control register (not shown) which stores a 15 value which can be set by the user of the voice communication device 1. In a first setting the control unit of the switch 4 uses a conventional learning algorithm to learn the MAC addresses of all devices which are connected to the switch 4 (including those devices which are connected to the switch 4 via the LAN and the port 11), and the association between those addresses and 20 the ports 9, 11, 14, 15, 17, 19.

In a second setting of the first control register, however, the learning algorithm is different in that switch 4 does not store the MAC addresses of the devices connected to the switch 4 via the port 11. In other words, the port 11 is a "first 25 port" as defined above. It continues, however, to learn the addresses of the devices connected to the other ports 9, 13, 15, 17, 19 (note however that since the port 9 is always connected to the same circuitry 7, it may not be necessary to learn the MAC address associated with this port; in other words, this MAC address may be preprogrammed).

Generally, a small number of devices will be attached to the switch 4 using the ports 13, 15, 17, 19 (e.g. typically one device per port), so the maximum number of data items (i.e. MAC addresses and their associations with ports) which the switch 4 is required to learn is typically not more than about 5, not 5 more than about 10, or not more than about 15.

10 The data switch may control a second control register (e.g. set by an external signal), which determines whether the state operates in a "learning state" or a "normal state".

15 In the learning phase, the learning of this data is via conventional learning techniques (as sketched above) based on the origin address of data packets received through the ports. For example, when a data packet is received from any of the ports having a destination address which is not associated in the address table with any of the ports, that data packet is transmitted to all the 15 (other) ports 9, 11, 13, 15, 17, 19.

20 However, once it is determined that the addresses of all the devices connected the ports 9, 13, 15, 17, 19 have been inserted into the address table, the second control register is reset to indicate that the switch should operate in the normal mode, and the operation of the switch 4 is as follows.

25 If the switch receives a data packet from any of the ports 9, 11, 13, 15, 17, 19 having a destination address which corresponds (according to the MA table) to one of the ports 9, 13, 15, 17, 19, then the switch transmits the data packets to that switch. In this way the switch makes possible communication between the circuitry 7 and the devices attached to the sockets 23, 25, 27, 29, and allows for incoming data packets from the LAN to be directed to the circuitry 7 or the correct one of the devices attached to the sockets 23, 25, 27, 30 29.

If a data packet is received from the port 11 which does not have a destination address associated (according to the MA table) with any of the ports 9, 13, 15, 17, 19, then the packet is discarded.

5

If, a data packet is received from any of the ports 9, 13, 15, 17, 19 which does not have a destination address associated (according to the MA table) with any of the ports 9, 13, 15, 17, 19, then the packet is transmitted only to the port 11, which transmits it in turn to the LAN. It is then the job of the LAN to 10 ensure that the packet reaches the device associated with the packet's destination address.

Thus, in comparison to the first setting of the first control register (and in comparison to the known systems), the second setting of the first control 15 register permits the memory usage of the switch to be very much reduced, while maintaining the same performance level.

Note that the first control register may be identical to the second control register. In other words, it may be the same control bit which in a first setting 20 indicates that the switch should operate as in the prior art, and in a second setting indicates that learning has been completed and that the switch should (i) no longer learn addresses associated with the port 11, and (ii) should discard all packets received from the port 11 which are not recognised by the table.

25

Note that in this case, by the time the control bit is changed over at the end of the learning phase, the address table may have already learnt some MAC addresses associated with the port 11. In this case, during the normal phase, data packets received from the ports 9, 13, 15, 17, 19 carrying a destination 30 address associated by the address table with the port 11 are forwarded to the

port 11, just as they would be if their destination addresses were not recognised at all by the address table.

Although only a single embodiment of the invention has been described,

- 5 many variations are possible within the scope of the invention as will be clear to a skilled reader. In particular, although the switch is shown in use in a voice communication device, other applications of the switch are possible, particularly in cases when the switch is connected to a network via a single one of its ports and through each of its other ports to a small number of
- 10 devices, such as exactly one other device, no more than 2 devices, or no more than 3 devices.

Claims

1. A data switch having a plurality of ingress/egress ports and for transmitting data packets including a destination address, the data switch having address table construction means for generating a table containing associations between ports of the switch and MAC addresses of any devices connected to the switch via those ports, the address table construction means being operable to construct said table in respect of all but a first one of the ports.
5
2. A data switch according to claim 1 in which the address table construction means is further operable to construct said table in respect of all of the ports, according to the setting of a control register.
10
3. A data switch for transmitting data packets including a destination address, the data switch comprising:
 - a first ingress/egress port and a plurality of second ingress/egress ports;
15
 - table storage means for storing a table containing associations between the second ports of the switch and MAC addresses of any devices connected to the switch via the second ports;
 - a switching fabric, and
20
- the control unit being arranged, upon receiving a data packet from any of the second ports having a destination address which is not stored in the table, to control the switching fabric to transmit the data packet to the first ingress/egress port.

4. A data switch according to any preceding claim which is connected to a communication network via the first port.

5. A switch according to any preceding claim having at least one port other than the first port arranged to receive and transmit voice signals.

5 6. A voice communication device for supporting communication with a second device, the communication device comprising:

a microphone,

a speaker,

10 circuitry for transforming sound signals received from the microphone into data packets and for transforming data packets into control signals for the speaker, and

a data switch according to any of claims 1 to 4 connected to the circuitry.

7. A voice communication device according to claim 6 further including 15 means for connecting one or more of the ports other than the first port to devices which each have a MAC address.

8. A voice communication device which is connected to a communications network via the first port of the switch.

9. A method of operating a data switch comprising a plurality of 20 ingress/egress ports, the method including:


generating a table containing associations between ports of the switch and MAC addresses of any devices connected to the switch via those ports,

the generation of the table including constructing said table in respect of all but a first one of the ports.

10. A method of operating a data switch for switching data packets including a destination address, the data switch comprising a first ingress/egress port, a plurality of second ingress/egress ports, and a memory storing a table containing associations between second ports of the switch
5 and MAC addresses of any devices connected to the switch via the second ports,

the method including, upon receiving a data packet from any of the second ports containing a destination address which is not stored in the table, transmitting the data packet to the first ingress/egress port.

FIG 1



INTERNATIONAL SEARCH REPORT

International Application No
PCT/02/00268

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 H04L12/44 H04L29/12 H04L12/64 H04L12/56 H04L12/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	GB 2 359 222 A (3COM CORP) 15 August 2001 (2001-08-15) page 3, line 24 -page 4, line 8 page 4, line 29 -page 5, line 26 -----	1-5,9,10
A	WO 01 08360 A (BROADCOM CORP ;KADAMBI SHIRI (US); AMBE SHEKHAR (US); KALKUNTE MOH) 1 February 2001 (2001-02-01) abstract page 30, line 7 - line 27 page 31, line 7 - line 30 -----	1-5,9,10
A	US 2002/054589 A1 (GAINER JAMES JOSEPH ET AL) 9 May 2002 (2002-05-09) paragraph '0016! paragraph '0028! - paragraph '0030! paragraph '0039! -----	1-10

 Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *&* document member of the same patent family

Date of the actual completion of the international search

Date of mailing of the international search report

12 June 2003

20/06/2003

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel: (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Ngao, Y.S.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/02/00268

Patent document cited in search report	Publication date	Patent family member(s)			Publication date
GB 2359222	A 15-08-2001	NONE			
WO 0108360	A 01-02-2001	AU 6198900 A	EP 1197045 A1	WO 0108360 A1	13-02-2001 17-04-2002 01-02-2001
		US 2002186705 A1			12-12-2002
		US 6430188 B1			06-08-2002
US 2002054589	A1 09-05-2002	US 6353609 B1	AU 4582199 A	BR 9911360 A	05-03-2002 05-01-2000 13-11-2001
		CA 2334927 A1	EP 1088431 A2	WO 9966759 A2	23-12-1999 04-04-2001 23-12-1999